IEEE/EG International Symposium on Volume Graphics (2010)
R. Westermann and G. Kindlmann (Editors)

Concurrent Volume Visualization of Real-Time fMRI

T. K. Nguyen[1], A. Eklund[2], H. Ohlsson[3], F. Hernell[1], P. Ljung[4], C. Forsell[1], M. Andersson[2], H. Knutsson[2], A. Ynnerman[1]

[1] Division for Visual Information Technology and Application (VITA), Linköping University, Sweden
[2] Division of Medical Informatics, Department of Biomedical Engineering, Linköping University, Sweden
[3] Division of Automatic Control, Department of Electrical Engineering, Linköping University, Sweden
[4] Siemens Corporate Research , Princeton, USA

Abstract

We present a novel approach to interactive and concurrent volume visualization of functional Magnetic Resonance Imaging (fMRI). While the patient is in the scanner, data is extracted in real-time using state-of-the-art signal processing techniques. The fMRI signal is treated as light emission when rendering a patient-specific high resolution reference MRI volume, obtained at the beginning of the experiment. As a result, the brain glows and emits light from active regions. The low resolution fMRI signal is thus effectively fused with the reference brain with the current transfer function settings yielding an effective focus and context visualization. The delay from a change in the fMRI signal to the visualization is approximately 2 seconds. The advantage of our method over standard 2D slice based methods is shown in a user study. We demonstrate our technique through experiments providing interactive visualization to the fMRI operator and also to the test subject in the scanner through a head mounted display.

Categories and Subject Descriptors (according to ACM CCS): 'I.3.3 [Computer Graphics]: Viewing algorithm— I.3.7 [Computer Graphics]: Color, shading, shadowing, and texture— I.4.10 [Image Processing and Computer Vision]: —Volumetric

1. Introduction

Functional Magnetic Resonance Imaging (fMRI) makes it possible to map out intensity and spatial location of brain activity. The estimation of brain activity is based on the measurement of the blood oxygen level dependent (BOLD) signal. As the development of both equipment and methods for fMRI progresses very fast, the pulse sequences used for fMRI have increased spatial and temporal resolution and there has been a shift from studies of static activation maps towards studies of the neuro-dynamical activation of brain regions during task performance. In addition, advances in signal processing are allowing handling of fMRI signals in real-time. This enables studies of the dynamics in the time-series data while the test-subject is still located in the fMRI scanner. Among the potential uses of real-time fMRI we find cognitive training and planning of stroke rehabilitation, assessment of drug treatments, and pre-operative planning. The primary current use of real-time fMRI systems are, however, in research on fMRI technology and methodology and they serve as enabling tools for neuroscience research.

To fully exploit the potential of real-time fMRI there is a pressing need for more visualization tools that not only visualize static activation patterns, but also effectively deal with the increasing spatial and temporal resolution. We thus have the following requirements on the fMRI visualization:

- High resolution subject/patient specific anatomical context data
- Fusion of the fMRI signal into contextual data in real-time to support concurrency with streaming of live data of increasing temporal resolution
- fMRI signal presentation that does not obscure the anatomical features
- Highly interactive visualization with frame rates that allow for smooth visual exploration

To meet these requirements we have developed direct volume rendering (DVR) methods tailored for fMRI time-sequence data. The methods are deployed in a demonstrator set-up for real-time fMRI. First the anatomical brain context is obtained from an MRI T1-weighted volume in a pre-processing stage. After being co-registered with the reference T1 volume, the fMRI signal is treated as a light source in the DVR pipeline, making the anatomical context data emissive depending on the signal location and strength. The result is that the brain *glows*, emitting light in active fMRI

DOI: 10.2312/VG/VG10/053-060

regions. The system supports several different shading models including the local ambient occlusion method described by Hernell *et al.* [HLY07]. The main contributions of the paper can be summarized as:

- Interactive DVR of the T1-weighted volume with the fMRI signal used to steer emission
- Inclusion of local ambient occlusion shading to improve the visual clarity of the anatomical context and the emission based on the fMRI activity
- A user study demonstrating the effectiveness of the proposed approach compared to standard 2D representations
- Demonstration of the use of the processing and visualization pipeline through a set of real-time fMRI experiments

2. Related Work

Previous research on volume visualization of fMRI has, to a large extent, focused on the use of multi-volume rendering approaches. Stokking *et al.* [SZV01] describe a generic method, called normal fusion, to visualize the functional data within context of 3D surfaces extracted from the anatomical image data. First, the functional information is projected onto the extracted 3D surfaces independently of the viewpoint. Then a color-encoding scheme based on the HSV model is used to fuse the anatomical and functional information. Weiskopf *et al.* [WVE*03] used the software package 'Turbo-BrainVoyager', (Brain Innovation, Maastricht, The Netherlands) [Goe01], for processing and visualization of fMRI data in real-time. Activities are visualized by superimposing the analyzed functional data on anatomical slices. For 3D visualization, active regions corresponding to activities on the surfaces extracted from the anatomical volume are marked with different colors. Rößler *et al.* [RTF*06] present a slice-based multi-volume rendering technique, which intermixes slices of volumes by depth sorting. The authors also present advanced cropping techniques and render modes to highlight activities inside the brain. Schafhitzel *et al.* [SRWE07] propose a combination of semi-transparent isosurfaces (SSD) with DVR to visualize the functional data. Particularly, DVR is used to render activities inside the brain while SSD is used to render the surrounding brain structure. The use of line integral convolution to provide a sparse and informative representation of the anatomical cues, and hence improve the perception of the isosurface's shape, is also proposed. In these approaches, the structure of the surface is well conveyed with the use of diffuse shading. Jainek *et al.* [JBB*08] use 'vicinity shading' to improve the depth and spatial perception in volume rendered images of functional data. In addition to variety in rendering methods including non-photorealistic rendering and hybrid methods, a set of novel interactive tools for investigating the visualization is also presented. Beyer *et al.* [BHWB07] introduce a multi-volume rendering workflow, in which high-quality visualization with interactive frame rates is achieved through a GPU-based implementation. This workflow is extended by Reider *et al.* [RRRhOP08] with a proposed solution for overlapping multiple volumes. Joshi *et al.* [JSV*08] introduced an integration of interactive cropping tools into a multi-volume rendering framework, for exploration of multimodal data in 3D. In their work, irregular cropping techniques can be applied to both fMRI and anatomical data to produce useful visual representations. Recently, Firdaus *et al.* [FBR*09] propose a method to visually analyze the time dimension of the functional data. While an activation dissimilarity metric is proposed to capture the context of the neuro-functional phenomena being investigated, a hierarchical clustering procedure is developed to provide suggestions for interactive volume of interest selection from users.

Although standard fMRI examinations have been used for almost two decades now (see Belliveau *et al.* [BKM*91]), the area of real-time fMRI is still relatively unexplored. Cox *et al.* [CJH94] were amongst the first to perform real-time fMRI analysis combined with on-line motion correction to compensate for head movements [CJ99]. Nakai *et al.* [NBM*06] used the general linear model (GLM) on a sliding window as a real-time activity measure. As an application, the use of real-time fMRI as an aid to cure chronic pain [dMG*05] has brought a lot of attention to the topic. The interested reader can find an overview of possible future applications of real-time fMRI in [deC08]. The recent developments of fMRI analysis have been focused on brain activity classification in real-time. Laconte *et al.* [LPH07] used a support vector machines (SVM) classifier to decide which of a number of stimuli could have caused a certain activity pattern. Eklund *et al.* [EOA*09] used a neural network for classification of brain activity to control a dynamical system.

3. A System for Real-Time fMRI

At the core of the system is a 1.5T Philips Achieva MRI scanner, which is a commonly used scanner for clinical MRI. The fMRI volumes acquired from a test subject are sent to a separate computer for real-time signal processing. The processed fMRI volume is passed to a second computer which generates the visualization. The visualization can be viewed by the operator, see fig. 1 for an example, and/or fed back and displayed to the test subject in the scanner using a head mounted display developed to work inside of the MRI scanner. The operator can interact with the visualization and control what he/she sees and what the test subject sees in real-time. The signal processing and visualization computer receive a new fMRI volume every other second while the visualization is updated more than 10 times per second. All units communicate over a local area network.

4. Signal Processing

Estimation of brain activity from real-time fMRI measurements is a challenging task. High noise levels and a weak BOLD signal make it hard to separate active from inactive regions. We here chose to work with periodic paradigms. This simplifies analysis since we only need to look for periodic fMRI signals but restricts the subject to periodically

Figure 1: An example of the view of the operator. The operator can easily interact with the visualization in real-time by using the keyboard and the mouse.

alternate between doing a task and rest (20s + 20s used throughout our experiments).

We use canonical correlation analysis (CCA) to detect periodically varying signals. CCA finds the two linear combinations of signals (one linear combination from one set and a second from another set of signals) which are as correlated as possible. Let w_x and w_y be the searched coefficients in the first and second linear combination and let $x(t)$ and $y(t)$ be column vectors containing the signals in the first and second set. The canonical correlation between $w_x^T x(t)$ and $w_y^T y(t)$ is then defined as

$$\frac{w_x^T C_{xy} w_y}{\sqrt{w_x^T C_{xx} w_x w_y^T C_{yy} w_y}} \tag{1}$$

where C_{xx}, C_{yy} and C_{xy} are the within and between sets covariance matrices, C_{xx} is e.g. defined by

$$C_{xx} = 1/N \sum_{t=1}^{N} x(t)x(t)^T. \tag{2}$$

The w_x and w_y that maximize the canonical correlation can be found by solving an eigenvalue problem, see [FBLK03] for details.

In our case, the first set contains a number of basis functions

$$x(t) = [\sin\omega t \ \sin 2\omega t \ \cos\omega t \ \cos 2\omega t]^T \tag{3}$$

where $\omega = 2\pi/T_p$, T_p being the period time of our paradigm (40s used here). By letting $y(t)$ be the fMRI measurement from a voxel, one could now obtain a measure of periodicity and hence activity in that voxel. The fMRI measurements are, however, contaminated by noise and $y(t)$ is therefore constructed by averaging together fMRI measurements from neighboring voxels. $y(t)$ will then consist of 5 elements and how their corresponding weights are defined is shown in fig. 2. The canonical correlation is computed voxel-by-voxel, as described above. In the computation of the within and between sets covariance matrices, C_{xx}, C_{yy} and C_{xy}, 80s of the most current fMRI measurements were used.

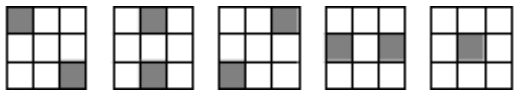

Figure 2: The 5 elements of $y(t)$ are constructed by weighing together fMRI measurements from nearby voxels as illustrated in the figure. In the 3-by-3 voxels arrangements, the voxel in the center is the one that the canonical correlation is computed for.

In order to compensate for head movement during the scanning, each fMRI volume is registered to the first collected fMRI volume. To do this we use a registration algorithm that is based on optical flow, but instead of the optical flow of the pixel intensity values that is normally used, we use optical flow of the local phase from quadrature filters. Phase-based volume registration is further described in [EAK10].

The registration between the low-resolution fMRI volume and the high resolution T1 volume is done by maximizing the mutual information between the volumes, as proposed in [VW95]. This is the standard approach for registration of volumes between different medical modalities. This registration is done before the real-time phase is started and takes about 1.5 minutes.

5. Contextualized fMRI visualization using volumetric illumination

The choice of illumination model used in volume rendering has been shown to have a major impact on the spatial comprehension [LBB04]. In particular, shadows serve as an important depth cue [SSMK05]. The Local Ambient Occlusion technique, presented by Hernell *et al.* [HLY07, HLY09] is used in this work to achieve smooth shadowing effects, making the brain tissue structure easy to perceive. Furthermore, a novel approach is proposed where the fMRI signal is treated as a light source – mapping the intensity to the strength of the fMRI signal, and thus illuminating the anatomical T1 brain volume where brain activity is detected, examples of images showing the effect can be seen in fig. 4.

There are three major processing stages in the rendering pipeline used in this application, as illustrated in fig. 3. The pipeline is based on a Direct Volume Rendering (DVR) scheme, where the traditionally constant ambient term in the Blinn-Phong shading model is replaced with two components: an Ambient Light contribution, A^L, based on Local Ambient Occlusion (LAO) computed using the brain anatomy T1 volume, and an Ambient Emission contribution, A^E, also computed using LAO with the addition of the fMRI signal to drive the emission factor. Both A^L and A^E are computed and cached on the GPU in two separate volume textures. Since the anatomy is static, the Ambient Light volume only needs to be updated when the Transfer Function (TF) for the anatomy is changed. The Ambient Emission, however, will be updated when the fMRI signal changes, from

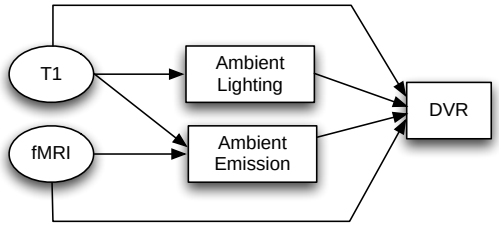

Figure 3: The volume rendering pipeline has three major processing blocks, Ambient Lighting, Ambient Emission and Direct Volume Rendering. The DVR stage will be executed for every update and change of view, whereas the Ambient Emission will be updated only when the fMRI signal is updated. The Ambient Lighting stage will be updated infrequently as it only depends on TF changes.

the real-time capture, or when the Emission Transfer Function is changed.

5.1. Local Ambient Occlusion

LAO is a local approximation of ambient occlusion which considers only voxels in a neighborhood around each voxel. Thus shadows and light emissions are captured from local features computing the incident light from a limited spherical region around each voxel, while ignoring distant features which have less impact. The resulting illumination volume is then used in a rendering pass as the ambient light term in the volume rendering integral. For the sake of clarity this section will briefly revisit the LAO principles.

The basis for LAO is the commonly used volume rendering integral, expressed by Max [Max95] as

$$I(D) = I_0 \cdot e^{-\int_0^D \tau(t)\mathrm{d}t} + \int_0^D g(s) \cdot e^{-\int_s^D \tau(t)\mathrm{d}t}\mathrm{d}s, \qquad (4)$$

where the first term represents the light coming from the background, I_0, attenuated with the optical depth, the integral of the extinction coefficient τ, and the second term represents the integration of attenuated light contributions, $g(s)$, for each location, s, along the ray. Traditionally, the light contribution function, $g(s)$, in eq. 4 has three components

$$g(s) = A(s) + k_d(L \cdot N)c(s) + k_s(N \cdot H)^p c(s) \qquad (5)$$

where $A(s)$ is the ambient light contribution, frequently simply a constant factor, as in $k_a(s)c(s)$, and $c(s)$ is the color classification of each sample. The second and the third term represent the diffuse lighting and the specular highlight, respectively, based on the half-angle technique by Blinn [Bli77].

The substance of LAO in the volume rendering pass is primarily integrated into the ambient term, $A(s)$ in eq. 5. Since $A(s)$ is typically independent of viewing parameters it is beneficial to compute this term separately and cache the result in a 3D texture that is only updated when the transfer function or volume data change. As mentioned above, A is furthermore broken up into an Ambient Lighting component, A^L, and an Ambient Emission component, A^E. The rest of this section will focus on the definition of A^L. The rationale for computing A^L and A^E separately will be detailed in the next section.

LAO captures the effect of a local volumetric ambient light source surrounding each voxel in the volume, attenuated by classified voxels in the neighborhood, accounting for the continuous range between transparent and opaque voxels. This volumetric approach, with light contribution to each g, creates a soft shadow penumbra, in contrast to only considering light at the boundary, that is, all light contribution given by I_0 alone. In addition, the volumetric light also achieves smooth results with less rays cast around each voxel, and thus reducing computational requirements.

In LAO the raycasting integral, eq. 4, is deployed over a set of K uniformly distributed rays originating at a location, x, in the volume and sum the contributions into $A^L(x)$. Thus capturing the total incident light as

$$A^L(x) = \frac{1}{K}\sum_{k=1}^{K} A_k^L(x) \qquad (6)$$

where the ambient light contribution, A_k^L, for each ray is defined by a reformulated version of the volume rendering integral, eq. 4. The integration direction is reversed, from center point to the boundary, as follows:

$$A_k^L(x) = \int_a^{R_\Omega} g_{A^L}(s) \cdot e^{-\int_a^s \tau(t)\mathrm{d}t}\mathrm{d}s \qquad (7)$$

where R_Ω defines the radius of the local neighborhood and a is an initial offset to avoid self-occlusion, τ is the optical density, defined by a transfer function for the T1 anatomy volume. To differentiate the light contribution at each sample location along the ray in the LAO processing, g_{A^L} is used instead of g.

A basic definition of g_{A^L}, using a volumetric light source spreading the light contribution evenly along the ray, is expressed by

$$g_{A^L}(s) = \frac{1}{R_\Omega - a}. \qquad (8)$$

So far, the Ambient Lighting processing stage, producing A^L, in fig. 3 has been described. In the left image of fig. 4 an example of rendering of the T1 brain context using A^L together with a small diffuse addition is shown. We now proceed to describe the processing for Ambient Emission, A^E.

5.2. Using the fMRI Signal as Illumination Source

One of the major ideas in this paper is to use the fMRI signal as a source for illumination. This is achieved by mapping the 3D fMRI intensity signal, indicating activity in the brain, through a separate transfer function defining a color emission, $c_E(s)$, for a location s along the ray. Adding color emission in the DVR stage requires a change to $A(s)$ used in

eq. 5. In addition the Ambient Lighting from the LAO stage, described in previous section, will be included such that

$$A(s) = A^L(s)c(s) + c_E(s), \tag{9}$$

with the ambient light, $A^L(s)$, being reflected from the classified T1 sample, $c(s)$. As can be seen in the left image of fig. 5 this produces an emissive effect but the fMRI signal is not yet illuminating the surrounding brain tissue.

To achieve an illuminating effect of the surrounding tissue the LAO algorithm is required to compute a first order light scattering event, based on the emission $c_E(s)$ from the fMRI signal. This information is computed and cached in a separate volume to hold A^E. Again the LAO algorithm is applied, but with $g_{A^E}(s)$ as in

$$A_k^E(x) = \int_a^{R_\Omega} g_{A^E}(s) \cdot e^{-\int_a^s \tau(t)\mathrm{d}t}\mathrm{d}s, \tag{10}$$

where

$$g_{A^E}(s) = c_E(s). \tag{11}$$

The sum of the rays for Ambient Emission generate A^E in a similar way as for A^L in eq. 6. It is also evident that A^L and A^E can be computed separately. Since $g_A = g_{A^L} + g_{A^E}$, the sum of the ray integrals over g_A is equivalent to the total sum of the integrals over g_{A^L} and g_{A^E}, respectively.

The full expression for $A(s)$ used in the volume rendering integral to include ambient lighting, $A^L(s)$, and emission, $A^E(s)$, is thus given by

$$A(s) = \left(A^L(s) + A^E(s)\right) \cdot c(s) + c_E(s). \tag{12}$$

This ambient color contribution, invariant to light and view direction, is thus a combination of both volumetric light and emitted color from the fMRI signal, reflected against the classified T1 sample's color $c(s)$, the traditional post-classified associated color. The direct contribution of fMRI signal emission, $c_E(s)$, is also included in the DVR process. In fig. 4 middle and right the resulting images for two different activation patterns are shown.

5.3. Visibility Enhancement

An additional feature has been added in order to ensure that signals deeper inside the brain will become visible while the visual context cues are maintained. For the samples in the region containing the fMRI data, the opacity in the ray casting pass, embedded in $c(s)$, is lowered if there is no signal emission at the given voxel location. The adjusted opacity, α', is computed as

$$\alpha' = \mathrm{clamp}(\alpha \cdot \mathrm{adjust\_threshold}, \alpha_{\min}, 1.0)$$

The opacity reduction is limited by a lower threshold $\alpha_{\min}$, essentially the transparency factor applied to all non-emissive voxels. The function clamp is used to keep the value within the range $[\alpha_{\min}, 1.0]$. The result is shown in fig. 6.

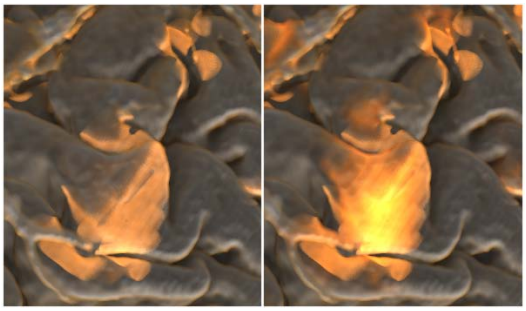

Figure 5: In the left image emission is only used in the raycasting pass. In the right image emission is applied both in the LAO and the raycasting passes, resulting in an increased glow effect around active regions. It is noticeable that the emission from within the brain "shines through" when LAO emission is used.

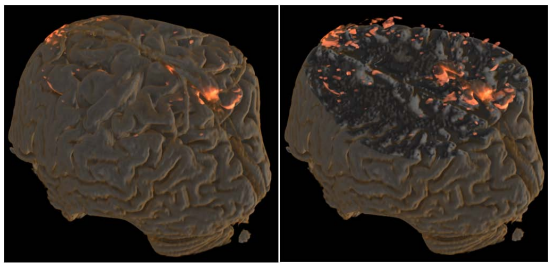

Figure 6: In the right image voxels without fMRI activity have their opacity reduced by a factor of $\alpha_{\min}$. In this image, the factor is set to 0.02. The adjustment makes interior signals more visible but still provides context in active areas.

Another mean to investigate the signals is by applying clip planes to the context and the signal. This makes it possible to better study the location and penetration of signal patterns. In fig. 7 the use of clip planes is illustrated.

6. Evaluation of method

To investigate our proposed fMRI DVR approach and to compare it with traditional 2D slice visualization we conducted an evaluation with 4 participants. The objective was to assess the users' ability to perceive spatial information of active regions (position) and also the intensity of the activity.

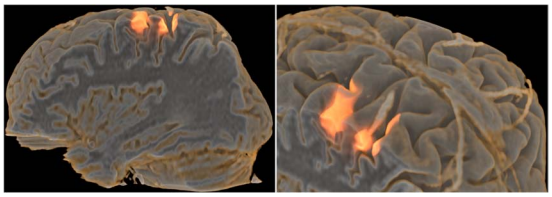

Figure 7: The use of clip planes to study the interior extent of the fMRI signal. In this case the signal represents activation of the motor cortex from movements of the left hand.

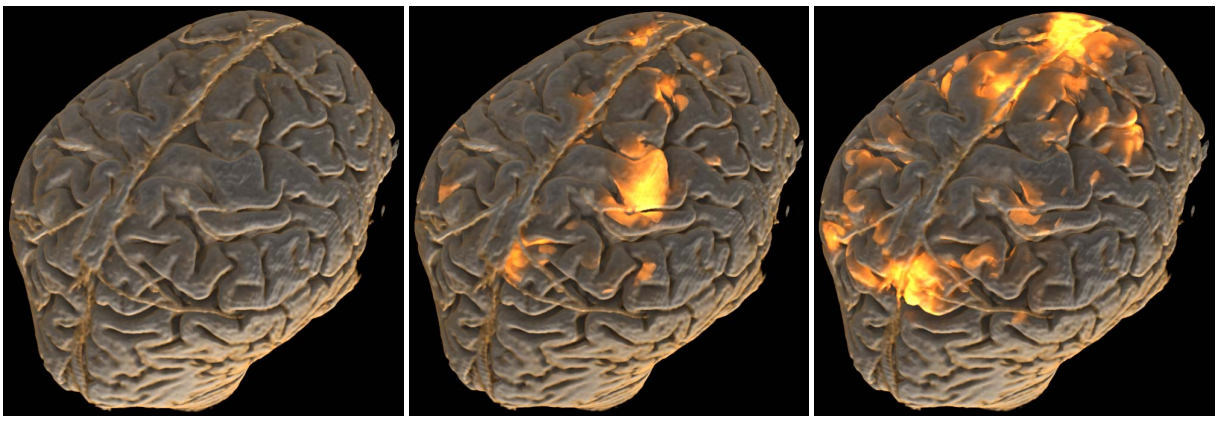

Figure 4: *Left:* The anatomy is rendered using the local ambient occlusion shading model, the A^L reflective light, which enhances the perception of depth. A diffuse shading component is also applied. *Middle:* The fMRI signal rendered using the LAO emission, i.e. including A^E. The image represents the brain activation during repeated motion of the left foot. *Right:* The activation during mathematical problem solving. The instructions, as well as the visualization of the brain activity as shown in these images, are shown to the subject in a head mounted display, and thus there is significant activation of the visual cortex.

We presented two vertically aligned images on a computer screen. Each image showed a brain activity pattern and the task was to compare the images and judge if they showed the same activity (same position) or not. In the slice visualization each image consisted of 30 slices arranged in a 6×5 matrix. Each slice had the active region marked in red, see fig. 8. In the DVR each image showed side-view and top-view of a brain volume with the active region marked in red.

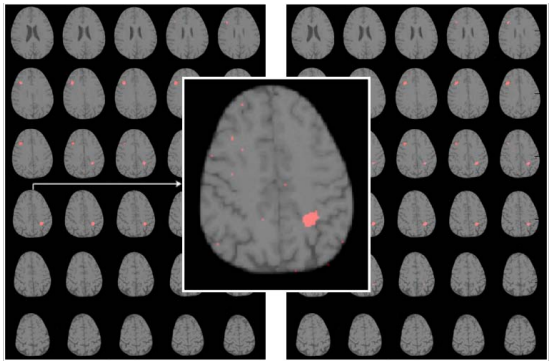

Figure 8: The 2D slice representations used in the evaluation. The task is to tell if the set of slices on the left and the right represent that the same activity or not. The inset shows a close up of one of the slices indicating the actual screen resolution of the experiment.

We used two synthetic data sets, A and B, to generate the brain activities. In data set A two activity patterns differ by a translation of 4 pixels in the x and y direction and 1 pixel in the z direction. In data set B, a translation of 10, 10 and 1.5 pixels was used. Also, random noise in position and shape of activity was introduced to simulate brain activity changes.

Table 1: The response data for all participants (P) for both visualization and both data sets. The figures within brackets show accurate/error/don't know responses respectively.

Test		P1	P2	P3	P4
Slice	**A**	(15/9/1)	(14/10/1)	(17/7/1)	(18/5/2)
	B	(24/0/1)	(23/1/1)	(22/3/0)	(23/2/0)
DVR	**A**	(20/5/0)	(19/2/4)	(21/4/0)	(24/1/0)
	B	(22/2/1)	(25/0/0)	(24/0/1)	(25/0/0)

The experiment was divided into two parts and half of the participants started with the slice visualization and the other half with the DVR visualization. In each part the participant viewed a randomized sequence of 50 pairs of images, 25 pairs from data set A and 25 from data set B. The presentation time was 10 s and they then responded to the question of whether the two images were: "same", "different" or "I don't know".

Table 1 shows the evaluation results. Performance was better for DVR than for the slice visualization for both data sets. As the displacement between two different activities got smaller and more difficult to perceive (data set A) performance improved by 31% using DVR compared to the slice visualization. This speaks in favor of our approach and future evaluations will expand on the results, providing a more comprehensive assessment of our approach.

7. Results

Experiments - The real-time fMRI visualization pipeline was tested using a 48 year old healthy male subject. The subject in the scanner was given instructions and notified of the onset and end of the periodic stimuli through a head mounted display. A green box indicated activity and a red box indicated rest. The period of the stimuli was 40 seconds, i.e. 20

Table 2: Signal Processing performance.

Signal processing step	Time (s)
fMRI Registration	0.8
Calculation of activity values	0.6
Total	1.4

seconds of activity and 20 seconds of rest. During the experiment the subject performed a number of different activities such as left and right hand activity, mathematical calculations, left and right foot activity and language tasks. In the experiments, the visualization was fed back to the subject in the scanner, as well as being displayed outside the scanner on a monitor. All data used to generate the images shown in this paper were obtained in this experiment.

Data acquisition and signal processing - Before the real-time phase could be started, a high resolution T1 weighted volume had to be captured. The resolution of the T1 volume was 140 slices, each with a resolution of 240 by 240 pixels. The T1 volume took about 5 minutes to acquire. The volume was then resampled using tricubic interpolation to a grid with 512×512×256 voxels. The physical size of a voxel is then 0.479×0.498×0.547 mm. Table 2 shows performance measurements of the described fMRI signal processing technique. The motion compensation and the calculation of activity is fast enough to collect a volume of data every other second. The resolution of the fMRI data is $64 \times 64 \times 22$ voxels, where each voxel has a resolution of $3.75 \times 3.75 \times 3.75$ mm. This means that we cover 8.3 cm of the brain in the sagittal direction. To obtain high BOLD contrast the echo time (TE) was set to 40 ms and the repetition time (TR) was set to 2000 ms.

Visualization - All of the volume rendered images have been generated using a standard PC equipped with an Nvidia GTX 285 graphics card with 1GB of graphics texture memory. The LAO map is computed with a, the initial offset to avoid self-occlusion, set to 0.8 voxel, the radius of spherical support, R_Ω, set to 2 blocks (1 block contains 16 voxels) and the number of steps per ray is 31 (see eq. 7).

Table 3 shows the performance for the calculation of the LAO map of the reference T1 volume and the fMRI signal. Initially, the LAO map of the reference T1 volume is calculated to provide a high resolution context with shadow effects. Since LAO requires processing of a neighborhood around each voxel in the volume, it is computationally demanding, as shown in the first column of Table 3, to compute the LAO for the full T1 brain volume. The fMRI signal is, however, only sampled in a subsection of the T1 brain and as shown in section 5, the emissive LAO from the fMRI signal can be computed separately. This makes it possible to calculate the LAO emission from the fMRI signal almost an order of magnitude faster then the T1 LAO computation. Furthermore, the signal is sparse and represents localized active regions. We can thus further improve the performance of

Table 3: Performance measurements for updating the LAO contribution of the T1 brain and the fMRI signal.

LAO calculation		T1 (s)	fMRI signal (s)	
			full calculation	optimized calculation
Initialization	32 rays	5.38		
	64 rays	10.98		
	128 rays	22.02		
fMRI signal change	32 rays		0.82	0.21
	64 rays		1.54	0.38
	128 rays		3.02	0.68

the fMRI LAO calculation by skipping voxels that are not illuminated by the signal. This yields an additional 4× speed up, as shown in the third column of Table 3. During the visualization, a change of transfer function can also be made. This requires both the recalculation of the LAO map of the reference T1 brain volume and the recalculation of the emission from the fMRI signal. In between TF changes and fMRI signal updates the DVR frame rate is 10.7 FPS. By ignoring the emission in the LAO stage the DVR frame rate can be increased, at the cost of visual quality, to 12.7 FPS. In this setting there is also no overhead associated with fMRI signal updates.

8. Conclusion and Future Work

In this paper we have described a system for real-time fMRI and we have shown how the emission in direct volume rendering can be used to concurrently visualize fMRI signals in the anatomical context of an MRI volume of a patient/test subject specific brain. The local ambient occlusion method for volume rendering yields visually appealing images of the brain anatomy and a clear fMRI signal presentation while rendering takes place at interactive speeds. One of the main advantages of the method is that it effectively fuses the anatomy with the fMRI signal of lower resolution, thus guaranteeing that the activity is only shown where there is brain tissue, as defined by the current transfer function settings.

The real-time visualization aspect of the system also has wider importance for normal fMRI use. The visualizations constitute a tool for many research situations as well as future clinical use. The in-vivo response of the system and high quality visualizations provide one of the important steps towards further utilization of the potential of fMRI. Among potential application areas of the real-time system we can mention cognitive training and planning of stroke rehabilitation, assessment of drug treatments and pre-operative planning. For instance, before performing brain tumor surgery, important areas of the brain are often mapped using a conventional fMRI experiment.

In future work we will further explore the possibility to feedback the visualizations to the test subject and see this as a part of a bio-feedback loop. We will also conduct more in-depth evaluations of the presented method. Of special in-

terest is to compare how our method depicts signal strength as well as the spatial position in comparison to other 3D visualization methods. We will also apply our visualizations to existing clinical fMRI routines and perform a detailed evaluation of the visualizations in the context of these routines. Another important area is to develop ways to visualize the uncertainty of the fMRI signal.

Acknowledgments

This work was supported by the Strategic Research Center MOVIII, funded by the Swedish Foundation for Strategic Research, SSF, the Swedish Research Council (SRC) on grant number 621-2008-4257, the Linnaeus Center CADICS. The Center for Medical Image Science and Visualization (CMIV) has provided the infrastructure for the fMRI experiments and medical expertise for evaluations.

References

[BHWB07] BEYER J., HADWIGER M., WOLFSBERGER S., BUHLER K.: High-quality multimodal volume rendering for preoperative planning of neurosurgical interventions. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (2007), 1696–1703. 2

[BKM*91] BELLIVEAU J., KENNEDY D., MCKINSTRY R., BUCHBINDER B., WEISSKOFF R., COHEN M., VEVEA J., BRADY T., ROSEN B.: Functional mapping of the human visual cortex by magnetic resonance imaging. *Science 254* (1991), 716–719. 2

[Bli77] BLINN J.: Models of light reflection for computer synthesized pictures. *ACM SIGGRAPH Computer Graphics 11*, 2 (1977), 192–198. 4

[CJ99] COX R. W., JESMANOWICZ A.: Real-time 3d image registration for functional MRI. *Magnetic Resonance in Medicine 42* (1999), 1014–1018. 2

[CJH94] COX R. W., JESMANOWICZ A., HYDE J. S.: Real-time functional magnetic resonance imaging. *Magnetic Resonance in Medicine 33* (1994), 230–236. 2

[deC08] DECHARMS R. C.: Applications of real-time fMRI. *Nature Reviews Neuroscience 9* (2008), 720–729. 2

[dMG*05] DECHARMS R. C., MAEDA F., GLOVER G. H., LUDLOW D., PAULY J. M., SONEJI D., GABRIELI J. D., MACKEY S. C.: Control over brain activation and pain learned by using real-time functional MRI. *PNAS 102* (2005), 18626–18631. 2

[EAK10] EKLUND A., ANDERSSON M., KNUTSSON H.: Phase based volume registration using CUDA. In *International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2010). 3

[EOA*09] EKLUND A., OHLSSON H., ANDERSSON M., RYDELL J., YNNERMAN A., KNUTSSON H.: Using real-time fmri to control a dynamical system by brain activity classification. In *Lecture notes in Computer Science, Proceedings of International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'09)* (2009), Springer, pp. 1000–1008. 2

[FBLK03] FRIMAN O., BORGA M., LUNDBERG P., KNUTSSON H.: Adaptive analysis of fMRI data. *NeuroImage 19* (2003), 837–845. 3

[FBR*09] FIRDAUS J., BOONTHANOME N., RAGHU M., HAN W. S., STEFFEN S., MICHAEL K., ISTVÁN A. M.: Visual analysis of brain activity from fmri data. *Eurographics/IEEE-VGTC Symposium on Visualization* (2009). 2

[Goe01] GOEBEL R.: Cortex-based real-time fMRI. *NeuroImage 13* (2001). 2

[HLY07] HERNELL F., LJUNG P., YNNERMAN A.: Efficient ambient and emissive tissue illumination using local occlusion in multiresolution volume rendering. *Eurographics/IEEE-VGTC Symposium on Volume Graphics* (2007). 2, 3

[HLY09] HERNELL F., LJUNG P., YNNERMAN A.: Local ambient occlusion in direct volume rendering. *Visualization and Computer Graphics, IEEE Transactions on PP*, 99 (2009), 1 –1. 3

[JBB*08] JAINEK W. M., BORN S., BARTZ D., STRASSER W., FISCHER J.: Illustrative hybrid visualization and exploration of anatomical and functional brain data. *Computer Graphics Forum 27* (2008), 855–862. 2

[JSV*08] JOSHI A., SCHEINOST D., VIVES K., SPENCER D., STAIB L., PAPADEMETRIS X.: Novel interaction techniques for neurosurgical planning and stereotactic navigation. *Visualization and Computer Graphics, IEEE Transactions on 14*, 6 (Nov.-Dec. 2008), 1587–1594. 2

[LBB04] LANGER M., BÜLTHOFF H., BÜLTHOFF I.: Depth discrimination from shading under diffuse lighting. *Proc. 7th Tübingen Perception Conference (TWK* (Jan 2004). 3

[LPH07] LACONTE S. M., PELTIER S. J., HU X. P.: Real-time fMRI using brain-state classification. *Human Brain Mapping 28* (2007), 1033–1044. 2

[Max95] MAX N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics 1*, 2 (June 1995), 99–108. 4

[NBM*06] NAKAI T., BAGARINAO E., MATSUO K., OHGAMI Y., KATO C.: Dynamic monitoring of brain activation under visual stimulation using fMRI-the advantage of real-time fMRI with sliding window GLM analysis. *Neuroscience Methods 157* (2006), 158–167. 2

[RRRhOP08] RIEDER C., RITTER F., RASPE M., HEINZ OTTO PEITGEN: Interactive visualization of multimodal volume data for neurosurgical tumor treatment. *Eurographics/IEEE-VGTC Symposium on Volume Graphics 27* (2008). 2

[RTF*06] ROSSLER F., TEJADA E., FANGMEIER T., ERTL T., KNAUFF M.: Gpu-based multi-volume rendering for the visualization of functional brain images. *Procceedings of SimVis 2006* (2006), 305–318. 2

[SRWE07] SCHAFHITZEL T., ROSSLER F., WEISKOPF D., ERTL T.: Simultaneous visualization of anatomical and functional 3d data by combining volume rendering and flow visualization. *Medical Imaging 2007: Visualization and Image-Guided Procedures* (2007). 2

[SSMK05] SATTLER M., SARLETTE R., MÜCKEN T., KLEIN R.: Exploitation of human shadow perception for fast shadow rendering. *APGV '05: Proceedings of the 2nd symposium on Applied perception in graphics and visualization* (2005), 131–134. 3

[SZV01] STOKKING R., ZUIDERVELD K. J., VIERGEVER M. A.: Integrated volume visualization of functional image data and anatomical surfaces using normal fusion. *Human Brain Mapping* (2001). 2

[VW95] VIOLA P., WELLS W.: Alignment by maximization of mutual information. *Proceedings of the 5th International Conference on Computer Vision* (1995), 16–23. 3

[WVE*03] WEISKOPF N., VEIT R., ERB M., MATHIAK K., GRODD W., GOEBEL R., BIRBAUMER N.: Physiological self-regulation of regional brain activity using real-time functional magnetic resonance imaging (fMRI): methodology and exemplary data. *NeuroImage 19* (2003), 577–586. 2